AD-787 092

DETERMINING OPTIMAL POLICIES FOR ECOSYSTEMS

George B. Dantzig

Stanford University

Prepared for:

Office Naval Research
National Science Foundation
Atomic Energy Commission
National Institutes of Health

August 1974

i

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

AD-787092

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 74-11	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): DETERMINING OPTIMAL POLICIES FOR ECOSYSTEMS	5. TYPE OF REPORT & PERIOD COVERED: Technical Report
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): George B. DANTZIG	8. CONTRACT OR GRANT NUMBER(s): N-00014-67-A-0112-0011
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Department of Operations Research, Stanford University, Stanford, CA 94305	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: NR-047-064
11. CONTROLLING OFFICE NAME AND ADDRESS: Operations Research Program Code 434, Office of Naval Research, Arlington, Virginia 22217	12. REPORT DATE: August 1974
	13. NUMBER OF PAGES: 17
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report): UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

This document has been approved for public release and sale; its distribution is unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

SYSTEM OPTIMIZATION
OPTIMAL CONTROL
FIR FOREST
MARKOV MODEL
ECOLOGY
INSECT DEVASTATION
SIMULATION

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

SEE REVERSE SIDE

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. ABSTRACT

DETERMINING OPTIMAL POLICIES FOR ECOSYSTEMS

BY

George B. Dantzig

This paper is a review of research done primarily at IIASA. The problem of finding the optimal policy for controlling the spruce budworm -- an insect whose outbreaks from time to time do great damage to the fir forests of New Brunswick, Canada -- represents a rare opportunity to develop and to successfully apply the methodology of optimization. The two interacting populations, the tree and the insect, constitute about the simplest ecosystem of practical importance. A very detailed computer "simulation" model is used to evaluate and to compare proposed policies regarding when to apply insecticides and when to cut down trees. The model is considered by biologists to be sufficiently representative that its simulation on the computer can be viewed as one way to bring the real world into the "laboratory". The effectiveness of different policies can then be determined by trying them out on the simulation model.

In this paper we discuss how the simulator can be supplemented with optimization methods to determine an optimal policy, in particular how a Markov model is appropriate. Carlos Winkler, a graduate student at Stanford has developed a computer program and has obtained optimal policies. Efforts are now under way to compare (on the simulator) his results with those based on intuition developed by experts. There is promise that the optimizer's policies (although based on a simplified model which did not consider cross effects of disease spread) are superior.

TR 74-11

DETERMINING OPTIMAL POLICIES FOR ECOSYSTEMS

by

George B. Dantzig

TECHNICAL REPORT 74-11

August 1974

DEPARTMENT OF OPERATIONS RESEARCH

Stanford University
Stanford, California

AND

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
(IIASA)
Laxenburg, Austria

Research and reproduction of this report were partially supported by: National Science Foundation, Grant GP 31393X2; Office of Naval Research under Contract N-00014-67-A-0112-0011; U.S. Atomic Energy Commission Contract AT(04-3)-326-PA#18; and National Institutes of Health Grant GM-14789-05. This research was primarily supported by IIASA.

DETERMINING OPTIMAL POLICIES FOR ECOSYSTEMS*

by

George B. Dantzig
INTERNATIONAL INSITUTE FOR APPLIED SYSTEMS ANALYSIS (IIASA)
and
STANFORD UNIVERSITY

The problem of finding the optimal policy for controlling the spruce budworm -- an insect whose outbreaks from time to time does great damage to the fir forests of New Brunswick, Canada -- represents a rare opportunity to develop and to successfully apply the methodology of optimization.

Most ecosystems are characterized by the complexity of the many species interacting with each other, the land, water, and climate in which they live. Many of the quantitative relationships which exist between the various species and their environment are still not known. For this reason ecologists rightfully fear that monkeying with the ecosystem could upset some delicate balance. They worry that some minor species could suddenly become dominant causing the entire ecosystem to move towards some ecological disaster. It is easy to understand why ecologists generally react negatively to any proposal that might affect the ecosystem.

Ecosystem models that have been developed are descriptive of only those parts of the system that are well understood. They are, at best, only partial models. The forest management problem, however,

---

*Details of the method sketched here will be found in the paper by Carlos Winkler.

represents an exception. The two interacting populations the tree and the insect constitute (according to C. S. Holling) about the simplist ecosystem of practical importance. Much is known about the growth of the fir forest under a variety of weather conditions and insect infestation. This has permitted the development at the University of British Columbia of a very detailed computer "simulation" model. Brought to IIASA for further studies, it is used to evaluate and to compare proposed policies regarding the cutting down of trees and the application of insecticides. Later we shall discuss how the simulator can be supplemented with other methods to determine optimal policy but first we will say more about simulation itself.

The scientific approach is to analyze a complex system in order to find laws that relate the parts of a system to one another and then to synthesize the relations to predict the observed characteristics of the entire system. For the latter purpose computers can play an important role. They can be used to compute changes in the state of a modeled system over time. Often continuous time is replaced by discrete time jumps and continuous relations are replaced by discrete approximations. When the computer is programmed to make the calculations by discrete time jumps, it is often referred to as a <u>simulator</u>.

The budworm simulator is made up of a great many computer subroutines and statements based on the equations, tables and graphs used by biologists to relate the number of trees and insects that appear in one year to those that appear in the next. These relations depend on weather conditions, age and state of health (stress) of the trees, the number of trees cut down, and the quantity of insecticides sprayed. The collected set of these relationships is referred to as the <u>model</u>.

The model for the spruce-budworm fir tree situation is considered by biologists to be sufficiently representative of the real situation that its simulation on the computer can be viewed as one way to bring the real world into the "laboratory." The effectiveness of different policies can then be determined by trying them out on the simulation model. Given any starting status of the forest and given the ground rules (policy) as to when to cut or spray (and how much) -- a medium-sized computer can, at low cost, calculate the change in status of the forest model for the next 100 years (say). It can do this again and again under a variety of weather patterns and record the number of times the forest (as modeled) was devastated by the insects, the amount of lumber obtained, and the final state of the forest. The computer can be programmed to output summaries in the form of graphical displays and charts.

The model can also be validated by having the computer run through a number of years of past history and the predictions based on the model can be compared with the actual observations (of the New Brunswick forest). Such validations constitute an important way to gain confidence that the model is representative -- if not then it could lead to further refinements.

Many of the relations of a model may be stochastic. For example, given the density of eggs laid on a leaf and the weather, it may not be possible to predict exactly how many eggs will survive to become adults because of the host of other factors which effect egg survival but which have not yet been measured, analyzed and put into the model. Most of the relations used in the budworm model are deterministic

by this is meant that the expected value of each outcome is used instead of a distribution of possible outcomes each with its own probability of occurrance. The effect of weather was treated, however, stochastically. Without too much difficulty the simulator could have been programmed to accept stochastic relations throughout. The resulting model would then be quite sophisticated. It would take time, and work, however, to make the budworm model fully stochastic and therefore it is recommended this extension of the model should be considered only after senstitivity analysis of the present model makes it evident that it would result in substantially better policies.

Objectives

One place where the formulation of models runs into difficulty is in the identification and quantification of objectives. This is certainly true of the forest management problem where a variety of general objectives can be stated such as: (1) obtaining high yields of lumber, (2) preserving the water shed, (3) preserving the forest as a recreational area, and (4) making the forest resiliant to disease and draught. Sometimes objectives are enunicated only after a solution of a model is presented and its bad characteristics noted. Sometimes objectives are stated in the form of requirements. Thus the requirement to meet a certain standard of water purity may be viewed as a way to achieve an environmental objective.

In general there appears to be no completely satisfactory way to state what is to be optimized when a model has multiple objectives. Probably the best approach to an understanding of objectives is through

dialog. Any particular run of the model becomes a dialog at the time its solution is reviewed. New runs can then be re-initiated and the dialog process continued until a satisfactory compromise solution is obtained.

It is interesting to note that the simulator as developed by the British Columbia group had no explicitly stated objectives. Of course the specific policies that they built into their simulator were put there to bring about, if possible, certain "desirable" goals but these were not explicitly stated. When work began at IIASA to find "optimal" policies, it opened up the whole question of how to make the objectives more explicit and quantifiable. This research on goals has already resulted in two important side studies.* It is thus interesting to note that the forest-management model has also provided an unusual opportunity for developing methodology for goal definition, measurement, and tradeoffs -- one which could be applied to many other practical problems.

The simulator (as we have already noted) is a good way to compare various policies for cutting and spraying to see which among them is "best". This requires some explanation because the decision as to what criteria to use for comparing the effects of applying one policy with those of another, is far from settled. The simplest (and most obvious) single quantity to use for comparison is the discounted value of all timber harvested in the future. If it turns out that the

*Reference here is to the work of Bell and Clark.

policy that yields the maximum discounted value does not cause the forest to be in an undesirable state at some future time, then this purely economic criterion is likely to become the one accepted for comparison of policies. Should it result, however, in a solution in which (say) the lumber industry cuts down a large number of trees one year and a few the next, then such a solution may receive a lower rating than some other one that yields lower profits but has a more even employment pattern. It may also be given a low rating if it does not result in a good mix of stands of trees of various ages, since a mixed forest is desirable because it can support a greater variety of wildlife and because it is more attractive as a recreational area.

An objective that is considered most important to achieve is the following: A forest is said to be "robust" or "resilient" if none of a wide range of possible stresses can trigger the forest to move in one or more years to a "bad" state. One way to achieve this, it is believed, is to have a good mix of trees of different ages since some ages are less vulnerable than others to annihilation by some disease or by extremes of weather.

Finding Optimal Policies

A policy for the forest model is a statement whether or not (1) to cut dow a tree and replant, (2) to leave it alone until the following year, or (3) to spray (if the latter, then the amount of spray to be used on the insect and when, must also be stated). As noted the simulator is very efficient if one wishes to compare one

policy with another. It can also be programmed so that one can vary one (or more) of the parameters defining a policy to see if some change in them (e.g., the amount of insecticide) will lead to policy improvement.

Because of the simplicity of the fir-tree budworm ecosystem this crude procedure might possibly be used to search for the optimal policy. However, for more complex systems this type of search is hopelessly inefficient unless it can be combined with a more analytic approach. A natural question to ask is this: is there any hope that a practical analytic method can be devised for finding the optimal policy? Let us note some difficulties: The New Brunswick Forest consists of 265 sub-forests, each with its own distribution of trees and different ages from 0 to 60 years (or more) in various states of stress (health) and degree of infestation by the budworm. Let us note the complicated mathematical equations, tables, graphs, and stochastic weather factors which govern the change of the forest from one year to the next. One is naturally discouraged by all this complexity from trying the analytic approach unless one can find some way to simplify the model. If the resulting policy based on a simplified model turns out, to be better (when compared using the simulator), than one obtained using (say) intuitive rules of thumb then this would justify the use of more analytic tools.

Selection of an Analytic Model

The approach we have taken is to regard the simulator as a means of bringing the real world into the laboratory. The various policies (whether obtained by common sense, or by common practice or through the use of an "optimizer") can always be compared by making a sufficient number of runs on the simulator. An analyst weak in analytic skills, poorly trained in the formulation of models, poorly informed about algorithms for solving classes of models, or unfamiliar with software availability may well opt to run many cases on the simulator to see if local improvements in a proposed policy is possible. Most simulation efforts unfortunately end up this way. Unfortunate because the high cost of using simulators to test many cases usually exhausts the patience and funds of sponsors to support development of an optimizer. If these funds had been used instead to develop a simplified model, then the process of determining an optimal policy for the simplified model could serve as a "brain" for the simulator and would have resulted in significantly better policies being found.

Generally speaking there are two types of analytic models that have had many successful applications: (1) "linear programming", and (2) "dynamic programming" models.

The first, the linear programming model, is characterized mathematically by a system of linear inequalities. Many kinds of non-linear relations can be practically approximated by such systems which can be both dynamic and stochastic. Software is available for solving such systems at reasonable costs even when they involve thousands of inequalities and variables.

The second, the dynamic programming model, is characterized by a dynamic system that moves from any given state in time to the next without being effected by the past history of how it arrived at its given state. Many practical models can be cast in this form. In practice, however, applications are narrowly limited to those whose "state space" may be approximated by a low number of cases. In our research, however, we have pursued an alternative possibility -- one which allows the state space to be multidimensional and continuous in certain components. We were able to do this by finding a practical way to approximate the "pay off" for each state if one follows henceforth an optimal policy.

For the Budworm Optimizer we used a mathematical model closely related to the dynamic program -- the so called <u>Markov Process</u>. At each point in time t, the system is some state $A, B, C, \ldots$. If in state A it will move to state A or B or $C, \ldots$, at time $t+1$ with probabilities $P(A|A), P(B|A), P(C|A), \ldots$; similarly if in state B it will move to A or B or C at time $t+1$ with probabilities $P(A|B), P(B|B), P(C|B), \ldots$, etc.

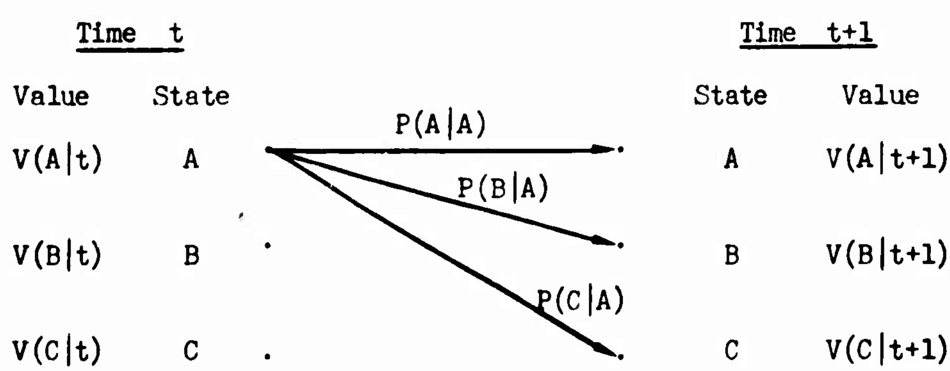

In our application these probabilities can be changed at a price by engaging in certain alternative actions. The problem is to find the best choice of these alternative actions. This is easy to do if we know the value $V(A|t+1)$, $V(B|t+1)$, ... of being in various states at time $t+1$. Thus the expected value $V(A|t)$ is given by

$$V(A|t) = P(A|A)\ V(A|t+1) + P(A|B)\ V(B|t+1) + P(A|C)\ V(C|t+1) \cdots .$$

If there are alternative actions in period t which can affect these probabilities, then the action that yields the maximum value of $V(A|t)$ is chosen. The procedure is thus a backward induction to time $t = 0$ but requires (in order to get it started) the knowledge of $V(A,t)$, $V(B,t)$, $V(C,t)$ for some future time $t = T$ in the future.

As noted a Markov type model was the one used for budworm study. The key idea used to develop this analytic model, was to view the single tree as an entity which changes state from year to year -- its state being defined by its age, stress, and the number of budworms it hosts. The tree, depending on the weather and whether or not it is sprayed or cut will (with certain probabilities) become one year older with certain stresses and egg densities or reverts to age zero and is replanted. If it were not for the spread from one timber stand to another of budworm eggs by the adult moth, this model has the merit that all other relations (equations, tables, graphs) can be used with little or no simplification or change. This leaves open the question of how to approximate the effect of egg contamination. We shall return to this important question later after we outline how the simplified model is solved.

For the simplified model we wish to find for every state (tree age, stress, and egg density) the optimal policy.

One way to determine optimal policy is to begin with a guess V_0 as to the entire discounted future value of a tree starting at age zero including the value of all its future harvesting and replanting (to time infinity) when we always carry out an optimal policy in the future with regard to the tree and its replantings. A tree planted a year from now, has present value of $.95\ V_0$ for its time stream from 1 year to infinity where 5% (say) is the discounted factor (without inflation). If for the moment we accept our guess V_0, we are in a position to evaluate the present value of <u>all</u> other states. One begins by noting that as far as harvesting the lumber of the tree now (or in the future) it does not pay to allow a tree to become older than 60 years (say). If so then the optimal policy is to cut it down and its present value $V_{60} = V_0 + L_{60}$ where L_{60} is the value of the 60 year old tree as lumber (less any cost for replanting it). To obtain the value V_{59} of a 59 year tree (which is in some state of stress and egg infestation) and, at the same time, to obtain the optimal policy, one compares the values obtained for each of the possible policies: (1) cutting it down, $V_0 + L_{59}$; (2) leaving it alone, $.95[pV_{60} + (1-p)V_0]$ where p is the probability of the tree living; and (3) spraying, $-S + .95[\bar{p}V_{60} + (1-\bar{p})V_0]$ where S is the cost of spraying and $\bar{p}$ is the probability of the tree living after it is sprayed. The policy which yields the highest value is selected as optimal. Note that the effect of random weather factors are part of the calculations (i.e., weather effects the probabilities of dying or the probabilities of moving from one state to another) so that values

(and optimal policies) of various states can be determined backwards from the highest age 60 down to age 0. If it turns out that our guess V_0 checks with the value V_0 obtained by the backward calculations we accept it -- if not then we revise our guess up or down until it does check.

Cross Effects of Contamination

So far we have treated the forest as a population of individual trees each undergoing its own private, independent transitions from one state to another. Unfortunately, the adult budworm (moths) lay their eggs over wide areas when there is little local foliage available. Thus the status of one area of the forest can effect other areas particularly those continguous to it and downwind. Now it may turn out that the optimal policies developed by the simplified model are such that at no time is there simultaneously a high density of eggs and a low amount of foliage in an area. In this event contamination effects will be low and the spread of eggs from one area to another will more or less balance the number of eggs it will receive from other areas. If so our study is complete.

If not, then the study of the best way to handle cross effects should be considered as part of the future research at IIASA. One approach that might be investigated is the following. First use the simplified model to provide approximate value functions at the end of the first year for each status. Next estimate the amount of contamination from other areas assuming that the other areas will use the policies of the simplified model. A region then decides its own policy

as the one which maximizes its own expected value when it includes as a cost (penalty) the cost it inflicts on the other regions by its contamination. It is further recommended that those regions be adjusted first which are more downwind and then those which are less so. I believe this adjustment method will converge quickly (if recycled) and will suffice for adjusting the policies of the simplified model to take into account the spread of insect infestation from one sub-region to another.

At this writing the steps necessary to bring this study to an orderly conclusion are now taking place. These are the comparison (on the simulator) of the policies developed by the optimizer with those developed by experts which they obtained using common sense, common practice and mature judgment. There is promise that the optimizer's policies are already superior to those of the experts and that these will become accepted as the new operating policies. If not then we can expect that the analytic models will be improved in some way. For example the objective used may be replaced by others that will make the forest more resilient or new procedures will be added to optimizers to adjust for contamination effects. Thus we can expect this process of dialog leading to improved models and solution techniques will continue until there is a consensus that optimal (or near optimal) operating policies have been attained.